Cisco XR 12000 Series
Service Separation Architecture Tests

Introduction

In April 2005, Cisco Systems released the XR 12000 Series routers bringing the functionality of IOS XR to the edge of next generation service provider networks. At the center of IOS XR is the Cisco Service Separation Architecture (SSA) which allows service providers to securely virtualize public and private services across the same platform by splitting a single router into distinct physical and logical routing domains. SSA enables the complete separation of network and system resources between each logical routing instance on the Cisco XR 12000 and allows additional dedicated memory and processing to eliminate control plane bottlenecks.

Test Highlights

- ➜ **The tests confirmed the Cisco XR 12000 operates with true separation between different entities of logical routers with zero overlap or forwarding between them.**
- ➜ **Software and hardware upgrades in one separated router had no performance effect on other logical routers.**
- ➜ **CPU and memory resources on the different logical routers' route processors proved to be independent of each other.**
- ➜ **The Cisco XR 12000 demonstrated 100 % IPv4 forwarding rate at 10 Gbps line rate and low latency in tests using realistic service scenarios of 3,200,000 fully meshed flows.**
- ➜ **The Cisco XR 12000 performed at 100 % forwarding rate and low latency for mixed IPv4 and IPv6 traffic with access control lists and logging for unauthorized traffic.**
- ➜ **3,000,000 BGP routes were advertised successfully resulting in zero loss forwarding performance and low latency for all routes.**

Venue & Test Equipment

Cisco XR 12000

Cisco commissioned the European Advanced Networking Test Center (EANTC) to verify the performance and functionality of the Service Separation Architecture. The tests were conducted at Cisco's test lab in San Jose, California in May 2005. EANTC test engineers conducted all tests and evaluated the Cisco XR 12000 behavior. The testing environment included Spirent's SmartBits load generator in combination with TeraRouting Tester software 4.50. The SmartBits 6000-C chassis was equipped with six Terametrics 10-Gigabit Ethernet XFP cards and 20x 1000Base-SX ports on four port Terametrics fiber and copper modules.

General Test Setup and Methodology

The tests used a Cisco 12416 chassis running IOS-XR software configured as four separate logical routers. Each logical router was equipped with two Performance Route Processor 2 (PRP-2) modules and one

Tested by EANTC 2005

EANTC extensively tested the Cisco XR 12000 for scalability, robustness and functionality of the Service Separation Architecture. The tests showed the expected performance values, confirming Cisco's claims for logical router performance and resource independence.

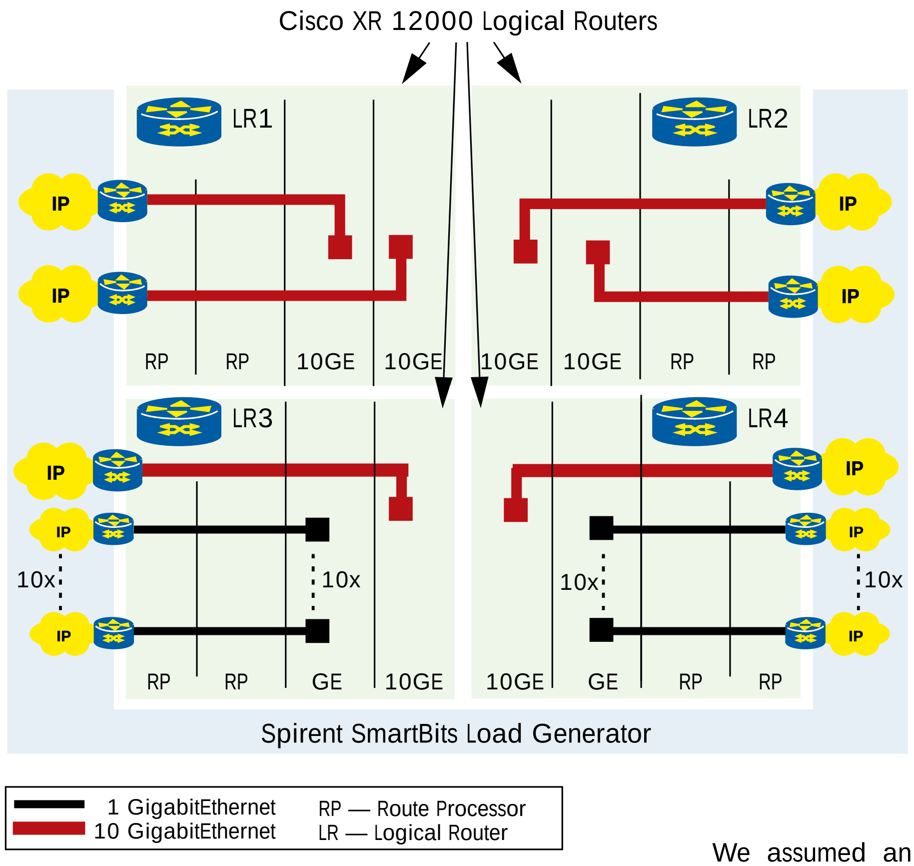

Logical Test Setup

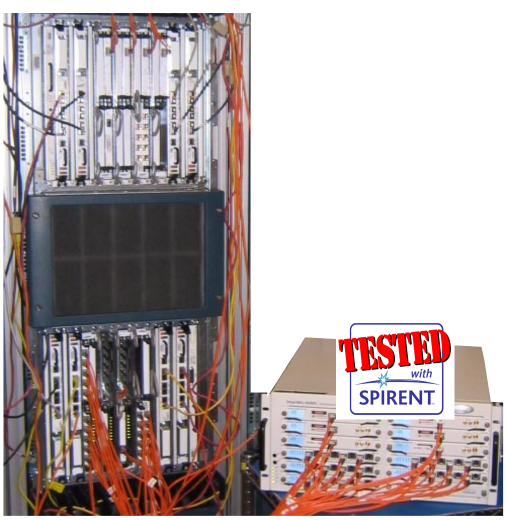

Physical Test Setup at Cisco's Test Lab, San Jose Cisco XR 12000 (left), Spirent SmartBits (right)

SPA Interface Processor 600 (SIP-600) module. Two logical routers were each equipped with two 1-port 10 Gigabit Ethernet Shared Port Adapters (SPA) with LR pluggable optics (SPA-1XTENGE-XFP; XFP-10GLR-OC192SR). The other two logical routers were each equipped with one 1-port 10 Gigabit Ethernet SPA with LR pluggable optics (SPA-1XTENGE-XFP; XFP-10GLR-OC192SR) and one 10-port 1000Base Ethernet SPA with SX pluggable optics (SPA-10X1GE; SFP-GE-S) to connect to the test equipment. The test setup is shown in the diagram above.

A 30-bit subnet mask was used for IPv4 transit networks, and a 124-bit network mask was used for IPv6. All tests without routing protocols and customer networks used flat networks (8-bit mask for IPv4 and 64-bit for IPv6).

We assumed an average bandwidth per emulated Internet flow of 50 kbit/s bidirectionally, resulting in 40,000 flows per GigE port and 400,000 flows per 10 GigE port.

IP Packet Stream Definitions

All load tests used a mix of packet sizes, representing Internet mix traffic (IMIX — see http://pma.nlanr.net/Datacube for a detailed explanation). IMIX traffic contains a realistic mixture of frame sizes that approximates real Internet traffic. IMIX packet sizes and distributions were defined as follows:

IPv4 Packet Mix (IMIX)		
Packet Size (Bytes)	Frame Size (Byte)	Bandwidth %
40	64	57 %
552	570	7 %
576	594	16 %
1500	1518	20 %

IPv6 Packet Mix		
Packet Size (Bytes)	Frame Size (Byte)	Bandwidth %
60	78	38 %
174	192	23 %
750	768	16 %
1500	1518	23 %

Using IMIX traffic allows the testing of the router under realistic conditions, as compared to single packet sizes tested sequentially.

1) Logical Router Separation Test

Test Highlights

➜ Different entities of logical routers configured in one XR 12000 chassis showed true separation. Traffic forced to cross router boundaries was not forwarded.

Test Objective

The aim of this test was to show the true separation between different entities of logical routers, with no overlap and or forwarding between them.

Test Methodology

This test setup used two logical routers (LR1 and LR2). Fully meshed IPv4 and IPv6 IMIX traffic streams were transmitted with wire rate across all ports of the configuration. Port1 and port2 belonged to logical router (LR) 1, port3 and port4 to LR2.

By inspecting the forwarded and lost packets, we determined that there was no packet forwarding between the separate logical routers.

One additional test run for an extended time (300 seconds) to check the configuration for long-term robustness.

Test Parameter Settings	
Traffic Pattern	IPv4 / IPv6 full mesh
Frame Size	64 / 78 Byte
Traffic Rate	100%
Test duration	60 s / 300 s (robustness)

Test Results

Stream	Expected Result	Measured Throughput
LR1 / port1 <-> LR1 / port2	Permitted	549,450,548
LR1 / port1 <-> LR2 / port3	Blocked	0
LR1 / port1 <-> LR2 / port4	Blocked	0
LR1 / port2 <-> LR2 / port3	Blocked	0
LR1 / port2 <-> LR2 / port4	Blocked	0
LR2 / port3 <-> LR2 / port4	Permitted	549,450,548

Results Graph

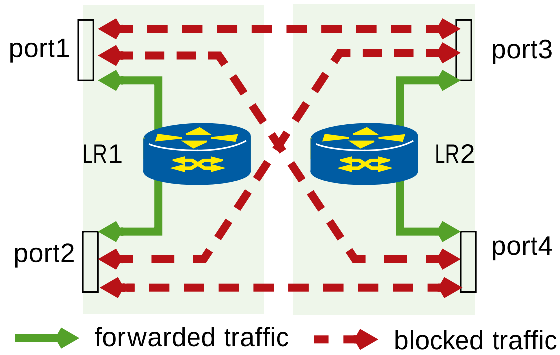

EANTC Test Analysis

As expected, we observed wire-speed forwarding for the flows within a logical router and 100 % packet loss for the traffic streams crossing the borders between logical routers. The system under test showed effective blocking of inter-logical router traffic.

2) Availability and Resource Independence

Test Highlights

➜ Software and hardware upgrades in one separated router had no performance effect on the other logical routers

➜ CPU and memory resources on the different logical routers route processors proved to be independent of each other

Part A: General Software Upgrade

Test Objective

The aim of this test was to observe the behavior of the router during software upgrades.

Test Methodology

Before initiating the software upgrade, we started a BGP test (see below, test case #6) to verify that the router maintained all peer adjacencies and continued to forward traffic during the software upgrade. In the next steps, the software of the so called *Owner Logical Router*, LR1, was upgraded (software upgrades on the owner LR affect all other LRs). We installed a *Package Installation Envelope* (PIE) including additional routing software. This upgrade also restarted the active BGP process as expected.

Test Results

Software Upgrade	Results
IPv4 / IPv6 Packet Rate	7,077,129 pps
Packet Loss	0
IPv4 / IPv6 Throughput, bidirectional	40 Gbit/s
IPv4 / IPv6 Latency	23.5 µs (microseconds)
BGP Adjacencies	Recovered during graceful restart interval

EANTC Test Analysis

As expected, the software upgrades/downgrades did not affect the forwarding and routing instances of the XR 12000. No packet loss or variation in forwarding latency was observed; latency did not increase indicating that there were no packets buffered; and the BGP sessions were interrupted for a short time as expected, but recovered within the graceful restart interval so that no routes were lost.

Part B: Hardware Module Exchange Procedure

Test Objectives

The aim of this test was to observe the behavior of the router during hardware module exchanges.

Test Methodology:

We reran another eBGP peering test. While the test traffic was transmitted, we

1. first removed the primary and then the secondary route processor of module of a non-owner LR (LR2)
2. removed and reinserted one unused SPA (shared port adapter) and verified traffic on the same LR,
3. removed and reinserted a switch fabric module with and without prior software shutdown,
4. removed and reinserted a line card (SIP) and observe traffic on other LRs.

The test streams were analyzed in regards to latency, throughput and frame loss.

Service Separation Architecture

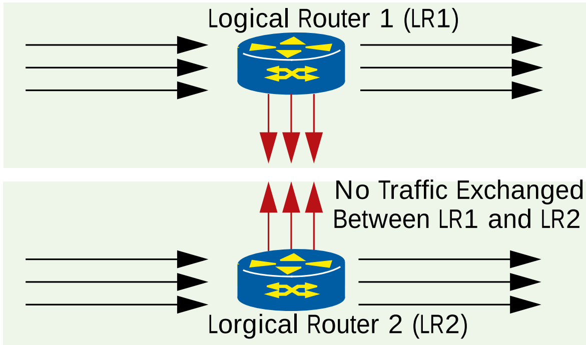

Test Results

1. The removal of the active RP had no affect on the forwarding parameters of LR2 or any other LR. Also the pulling of the second RP showed no packet loss or increased latency in any of the LRs (the test duration was not long enough to see the line cards rebooting after a certain delay)
2. The removal and reinsertion of an unused SPA worked without any influence on any of the other test streams. Initially, when we reinserted the SPA the line card holding that SPA rebooted. Cisco investigated the (software) issue and solved it.
3. The removal and reinsertion of one of the switch fabric modules, whether it was disabled by software before or not, did not show any packet loss or increased latency for any of the test streams. Some time after reinsertion, the Spirent SmartBits locked up due to BGP overload so that packet counters were invalid at the end of the test; Spirent investigated the issue at the time this report was printed.
4. The removal and reinsertion of a line card worked as expected: no packet loss or increased latency was observed as a result on influence of this action on the other logical routers.

EANTC Test Analysis

No unexpected packet loss or latency variation was observed when any route processor, interface module, interface processor, or switch fabric module were removed. Reinsertion of an interface module (SPA) caused a reboot of the interface processor which Cisco has resolved.

Part C: Logical Router Resource Independence

Test Objectives

The aim of these tests was to verify that the resources of one logical router are independent of the other logical routers.

Test Methodology

To verify that the CPU resources are independent, we started a cyclic SNMP walk on the logical router LR1 and ensured that the CPU load on LR1 increases significantly. Then we performed a eBGP peering test on all LRs to confirm that the forwarding performance is not affected. In parallel to the test, we monitored the CPU load of LR1 and LR2.

To verify that memory allocation processes and route processors are independent, we completely switched off a logical router during the BGP peering test (as opposed to pulling both route processors as in part B above). The test streams were analyzed in regards to latency, throughput and frame loss.

Test Results

Maximum CPU load of LR1 and LR2 during the test:

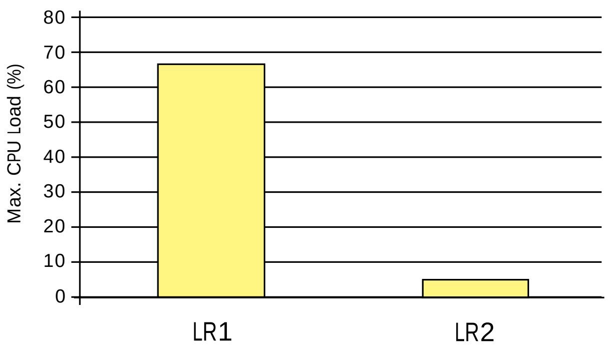

BGP Peering Test during SNMP walk on LR1	Results all LRs
IPv4 / IPv6 Packet Rate	7,077,129 pps
Packet Loss	0
IPv4 / IPv6 Throughput	40 Gbit/s
IPv4 / IPv6 Latency	23.5 µs

Switching off LR4 had no effect on packet loss or latency values of LR1, LR2 or LR3.

EANTC Test Analysis

The high CPU load of LR1 did not influence the performance of LR1 or any other LRs, thus verifying the independent CPU resources.

Completely switching off one logical router did not influence the traffic of other logical routers of the same chassis.

3) IPv4 IMIX Based Forwarding

Test Highlights
➜ Each separated router demonstrated full 10 Gbps wire speed IPv4 forwarding performance, for IMIX traffic as well as for traffic with homogeneous minimum (64 byte) frame size.

Test Objective

The objective of this test case was to determine the maximum IPv4 forwarding performance of the system under test under realistic conditions.

Test Methodology

The test used the IMIX IPv4 packet mix as previously defined. The test streams were generated in pairs for the two 10GigE ports of the first two logical routers and between the one 10GigE port and the ten 1GigE ports of LR3 and LR4. Within each of these streams, full-mesh flows were configured for different IP addresses to generate a total of 1,600,000 flows. No flows were configured between the separate logical routers.

With this set up we determined the maximum throughput and forwarding latency according to RFC 2544. One additional test was run at the maximum throughput for an extended time of 300 seconds to verify long-term forwarding robustness.

To find the IPv4 forwarding performance limits, we repeated the forwarding and latency tests with 64-byte Ethernet frames only.

Test Parameter Settings	
Traffic Pattern	IPv4 full mesh
Packet Size	IMIX and 64-byte only
Traffic Rate	100%
Test duration	60 seconds (packet loss) / 120 seconds (latency) / 300 seconds (robustness)

Test Results

IPv4	IMIX 60 s / 300 s	64-byte only
Packet Rate	72,950,342 pps	119,047,619 pps
Packet Loss	0	0
Throughput	40 Gbit/s	40 Gbit/s
Latency	24.5 µs	24.5 µs

EANTC Test Analysis

This test confirms Cisco XR 12000 10 Gbps wire speed forwarding with IMIX traffic as well as with small packets only traffic.

4) IPv4/IPv6 IMIX Based Forwarding

Test Objective

Test Highlights

➜ Each separated router demonstrated full 10 Gbps wire speed forwarding using IMAX IPv4/IPv6 traffic mix.

The objective of this test was to verify the maximum forwarding performance for mixed IPv4 and IPv6 traffic of the system under test.

Test Methodology

The test used the IMIX IPv4 and IPv6 packet mixes as previously defined. The flows were generated in the same way as in the IPv4 test case. Within each of the streams connecting the different ports of a logical router, fully meshed flows were configured for different IP addresses. IPv4 and IPv6 generated a total of 3,200,000 flows. As with the IPv4-only test, no flows were configured between the separate logical routers.

With this set up the maximum throughput and forwarding latency according to RFC 2544 was determined.

One additional test was run at the maximum throughput for an extended time of 300 seconds to verify long-term forwarding robustness.

Test Parameter Settings	
Traffic Pattern	IPv4/IPv6 full mesh
Packet Size	IMIX
Traffic Rate	100%
Test duration	60 seconds (packet loss) / 120 seconds (latency) / 300 seconds (robustness)

Test Results

IPv4 / IPv6	IMIX 60s / 300s
IPv4 / IPv6 Packet Rate	69,980,815 pps
Packet Loss	0
IPv4 / IPv6 Throughput	40 Gbit/s
IPv4 / IPv6 Latency	24.1 µs

EANTC Test Analysis

This test confirms Cisco XR 12000 wire speed performance with an IPv4 and IPv6 IMIX.

5) IPv4/IPv6 IMIX Based Forwarding with Services

Test Highlights

➜ Line speed forwarding was observed after adding security services like access control lists and logging of unauthorized traffic streams.

Test Objective

This test was designed to verify that the forwarding performance for mixed IPv4 and IPv6 traffic does not degrade when access control lists and logging for unauthorized traffic are configured.

Today, most IP switches in carrier and enterprise environments need to observe basic access control rules for security, traffic classification and SLA assurance, sometimes with basic intrusion detection mechanisms (ACL logging).

Test Methodology

For each protocol, IPv4 and IPv6, a 5,001 entry ACL was configured, where 5,000 entries were DENY and the last entry was a PERMIT-ALL. (ACL DENY criteria were pseudo random and not sequential, thereby preventing sequential ranges being converted into a single ACL entry. Also, 50% of the DENY entries were configured matching the IP addresses of the data streams we sent, but not matching the UDP ports of the data traffic. The ACLs were applied as both, ingress and egress filter on each port.

In contrast to the previous IPv4/IPv6 test case, UDP packets were sent rather than IP packets and UDP destination ports were varied between 5,000 and 7,500. We ensured that all traffic matched only the final "PERMIT-ALL" ACL entry, thereby forcing the switch to compare the traffic to all 5,001 ACL entries per port in both the ingress and egress directions (double lookups).

1. The initial test run was to determine the maximum throughput and average forwarding latency according to RFC 2544.

2. To confirm the ACLs were actively protecting the network, we re-ran the test with one traffic flow matching the 5,000th DENY entry in the ACL, the other IP addresses did not match any of the DENY entries. During this test run we monitored the system to verify that 100% of the traffic matching the DENY entry was dropped and that the remaining traffic was forwarded without loss.

3. Finally, we enabled ACL logging while sending the same traffic as in the second test run to observe if the ACL logging feature affected the performance or latency and that the details recorded could help identify the individual flows that were attempting to break through the ACL barrier.

To uncover potential long-term robustness issues we executed one additional test run at the maximum throughput for an extended time of 300 seconds.

Test Parameter Settings	
Traffic Pattern	IPv4/IPv6 full mesh with UDP
Packet Size	IMIX
Traffic Rate	100%
Test duration	60 s / 300 s (robustness)

Test Results

IPv4 / IPv6 With Services	Results
IPv4 / IPv6 Packet Rate	69,541,829 pps
Packet Loss	0
IPv4 / IPv6 Throughput	40 Gbit/s
IPv4 / IPv6 Latency	25.5 µs

No unexpected packet loss was observed. The test run with one flow matching the ACL only showed packet loss for the appropriate flow. No other flow was affected by packet loss.

Performance and latency values did not change when ACL logging was switched on. The following records were logged:

```
ipv4_acl_mgr[188]:%ACL-IPV4_ACL-6-IPACCESS-
LOGP: access-list IPV4_Test (50000) deny udp
140.0.0.1(1024) -> 141.100.0.1(5000), 1
packet
```

The extended duration test did not show any long-term robustness issues.

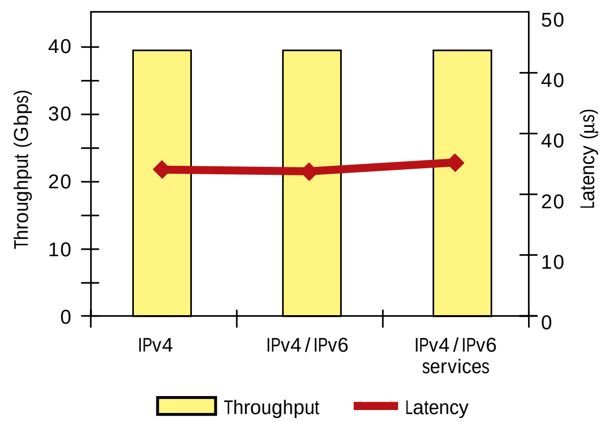

EANTC Test Analysis

The Cisco XR 12000 maintained line rate performance even with the addition of Access Control Lists and logging. Latency remained very small; it virtually did not increase compared to IPv4 forwarding without services. In both cases, the latency matched or exceeded expectations compared to informal industry standards.

6) BGP Peering Test

Test Highlights

➜ 3,000,000 BGP routes were advertised successfully resulting in zero loss forwarding performance and low latency for all routes

Test Objective

The goal of this test was to determine the performance of the BGP implementation using both IPv4 and IPv6 routes and traffic.

Test Methodology

This test emulated a realistic environment with a total number of 3,000,000 IPv4 and IPv6 routes distributed over all four logical routers. 20% of all routes were IPv6 routes. Traffic was sent over all routes using the an IMIX packet mix.

The following eBGP prefix distribution was used:

	Prefix Length	# of routes per logical router	distrib-ution per protocol
IPv4	/19	60,000	10%
	/24	360,000	60%
	/29	180,000	30%
IPv6	/48	15,000	10%
	/64	90,000	60%
	/96	45,000	30%

One peer for eBGP4 over IPv4 and one peer for eBGP4+ over IPv6 were configured per port. All advertised routes were verified with IPv4 and IPv6 IMIX traffic. One source IP address was used per route sending fully meshed traffic.

Test Parameter Settings	
Traffic Pattern	IPv4/IPv6 full mesh
Packet Size	IMIX
Traffic Rate	50%
Test duration	60 seconds

Test Results

BGP Peering	Results
IPv4 / IPv6 Packet Rate	7,321,190 pps
Packet Loss	0
IPv4 / IPv6 Throughput	40 Gbit/s
IPv4 / IPv6 Latency	22.9 µs

EANTC Test Analysis

All advertised BGP routes and updates were processed successfully, resulting in 100 % throughput for all routes during normal operation.

Conclusion

The Cisco XR 12000 fulfilled all of Cisco's performance and functional claims for all areas analyzed.

These tests re-affirm that the Cisco XR 12000 will be one of the most potent routers on the market and we believe that based on the results of these tests it will remain a popular choice for carriers planning next-generation networks.

About Cisco Systems

Cisco Systems, Inc. (Nasdaq:CSCO) is the worldwide leader in networking for the Internet. News and information are available at http://www.cisco.com.

About Spirent Communications

Spirent Communications is a worldwide provider of integrated performance analysis and service assurance systems for next-generation network technologies. Spirent's solutions enable customers to develop and deploy network equipment and services more economically and efficiently by emulating real-world conditions in the lab and assuring end-to-end performance of large-scale networks.

Spirent Communications is a wholly owned business group of Spirent plc, an international network technology company. Spirent, Spirent Communications and the Spirent logo are trademarks of Spirent plc.

http://www.spirentcom.com/

About EANTC

The European Advanced Networking Test Center (EANTC) offers vendor neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, ATM, VoIP, Triple Play, and IP applications.

EANTC AG
Einsteinufer 17, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/